Capturing Personal and Crowd Behavior with Wi-Fi Analytics

Utku Günay Acer[†], Geert Vanderhulst[†], Afra Mashhadi[†], Aidan Boran[†],
Claudio Forlivesi[†], Philipp M. Scholl[‡], Fahim Kawsar[†]
[†]Bell Labs, [‡]Albert-Ludwigs-Universität Freiburg

ABSTRACT

We present a solution for analysing crowds at events such as conferences where people have networking opportunities. Often, potential social relations go unexploited because no business cards were exchanged or we forget about interesting people we met earlier. We created a solution built on top of ubiquitous Wi-Fi signals that is able to create a memory of human trajectories and touch points. In this paper we elaborate on the technological assets we designed to perform crowd anlaytics. We present small wearable Wi-Fi badges that last for the duration of an event (up to 3 days) with a single charge, as well as network equipment that senses the signals radiating from these badges and contemporary mobile devices.

1. INTRODUCTION

Events such as conferences and summits are a catalyst for spreading ideas and transacting business. By providing a setting in which people from diverse backgrounds and with diverse purposes assemble, new social relations are established [12]. Events give rise to building trust and attracting new resources for firms [7] as well as identifying business and collaboration opportunities [4].

In this paper, we summarize a system that collects data to understand the impact of spatial structure on the behavioral dynamics of social groups assembled in an event. We report on an end-to-end system that collect data to accomplish these goals. We designed and developed a Wi-Fi solution comprising of wearable Wi-Fi badges and gateways. These gateways capture Wi-Fi signals radiating from wearable badges (the size of a credit card) and anonymous smart devices to create spatio-temporal trajectories of individuals discretely wearing them. By performing sensing at the network side rather than at the device end, we are able to keep mobile power consumption to a minimum. The system collects fine granularity data from the badges that are given to a selected group of people to achieve the first goal and coarse granularity data from the smart devices of the conference attendees to accomplish the second.

We have deployed this system in a large scale industrial event and capture the spatio-temporal trajectories of 2.5K+ attendees including two special groups: 34 investors and 27 entrepreneurs[1]. The event attracted 40K+ people from 134 different countries and took place over three days in early November 2015 in a European city. The $\pm$6000 sq. m. venue of the event was organized to create structural opportunities for the participants to meet and interact through technical presentation sessions, demo booths, marketing booths, startup pitch sessions as well as informal spaces.

WPA'16, June 25-30, 2016, Singapore, Singapore

DOI: http://dx.doi.org/10.1145/2935651.2935656

2. RELATED WORK

A rich body of literature has studied the crowd behavior during large events. One way to collect information is through a mobile app that uses a variety of resources such as GPS, Wi-Fi, Bluetooth, cellular, and other sensors from the mobile phone and integrate them to estimate the location of the user [3, 17]. However, installing and running the app on the critical mass of the event participants' devices is challenging. Hence, the majority of the previous work rely on opportunistic sensing.

Bluetooth tags are used by Jamil et al. to study the movement and community structure of pilgrimage during Hajj religious festival [11]. The ubiquity of mobile phones on the other hand made it possible to use network signals to trace and track people to study mobility and behavior without the need for specialized tags. Larsen et. al. use Bluetooth signals emitted from smartphones to analyze the behavior of participants to a large music festival [14]. Other works investigate the mobility and interactions of people by capturing Bluetooth signals from their discoverable devices with scanners [18]. Features about Bluetooth signals such as mean signal strength and the number of devices can also be leveraged to estimate the crowd density [22].

Prior work on crowd analysis widely used Bluetooth technology to collect information about participants. Though Wi-Fi signals have been widely leveraged for indoor localization [10, 15, 23], Wi-Fi based systems have not been deployed to understand the crowd behavior to the best of our knowledge. On the other hand, Wi-Fi based analytics methods have been used to detect face-to-face interactions [21], to track physical objects in time and space [2] and evaluating location reputation[16].

3. SYSTEM DESCRIPTION

In this section, we first discuss the design challenges we have faced in building the sensing system and present the components of the system.

[1]The consent to record the trajectory of those wearing the badges were given by the individuals themselves, and the consent to gather crowd analytics data from all the attendees were given by the event organisers.

3.1 Design Challanges and Decisions

Our system has two primary goals: (i) Retrieving the behavior of two groups of people, e.g. entrepreneurs and investors, and (ii) capturing the crowd dynamics in the event. For the first goal, guaranteed participation of a number of attendees from both groups were required. It was essential that these attendees produce signals at a high rate so that we can estimate their location at different points in time with fine granularity to construct their trajectories. With the second goal, we aim to establish the how the popularity of the prominent places in the event venue changes over time by looking at the spatio-temporal density. For such a feast, the system needs to collect data from all the attendees in the event, albeit with coarse granularity.

Past research has relied on technologies such as GPS [19] and Bluetooth [22, 11] as sensing modalities. Because the event was confined to indoors, GPS was excluded from our design alternatives. Though a Bluetooth-based approach offers accurate location estimates, it has a number shortcomings. Due to its short communication range, it requires a large number of scanners. It offers a limited participation with smart phones as the Bluetooth interface is turned off most of the times [6] and it requires a high energy cost [9].

In this study, we have relied on Wi-Fi signals from Wi-Fi enabled devices. This decision is grounded on findings that suggest people carry their smartphones with them most of the time [13] and their Wi-Fi is switched on particularly where/when free Wi-Fi service is available. Since Wi-Fi has a wider range than Bluetooth, it demands a fewer number of receivers capturing the wireless signals.

Wi-Fi and the underlying IEEE 802.11 standard [1] uses three types of link layer frames to facilitate communication the access points (AP) and non-AP devices. *Data frames* transports application layer traffic from higher layers. *Control frames* police devices accessing the wireless medium without causing an interference between each other. Every time a device has data to send to another device, they perform a handshake through Request-to-Send (RTS) and Clear-to-Send (CTS) frames. The other devices in the vicinity also listen to these frames and back off their transmissions during this exchange. *Management frames* are used to provide and maintain connectivity to the devices. It includes a number of subtypes such as authentication frames that facilitates securely connecting to a AP, association frames that bind the device to the AP, beacon frames for AP to announce its presence and its service set identifier (SSID), and probe requests to scan Wi-Fi networks in the vicinity. To receive a frame, the receiving device needs to listen the same Wi-Fi channel over which the transmitting device is sending the frame.

We have designed custom, isolated Wi-Fi gateways to capture Wi-Fi signals from attendees' smart devices. The gateways have no external connectivity to minimize interference and reduce congestion. Moreover, we have made two important design decisions:

- We capture only probe requests from the Wi-Fi enabled devices. Probe request - further referred to as *probe* - is a Wi-Fi management frame that is either directed to a specific network or broadcasted to any network by a mobile device while it scans Wi-Fi network .

- We use dedicated Wi-Fi badges that were distributed to a select number of investors and entrepreneurs. The badges emit probes systematically so that nearby Wi-Fi gateways can capture them, which are in turn used to estimate the positions of the wearers in the venues.

Though devices send data frames and control frames more frequently for their application traffic, we decided to only use probes. They are sent on the channel that the AP is operating on. However, with several APs present in the conference venue, each operating on a different channel to reduce interference from other devices, one cannot determine before hand what channel the gateway needs to monitor. On the other hand, when a Wi-Fi device scans the networks, it sends probes on every Wi-Fi channel. Therefore, even if a gateway monitors a single Wi-Fi channel, it receives probes from devices when they scan the networks. Given the crowded nature of the event and congested open Wi-Fi network, we expected that devices would often lose the connection, or look for a better connection resulting in the transmission of probes. Even if the device keeps the connection to a particular device, it still scans Wi-Fi networks with scanning interval defined by the vendor that allows device to be detected by the deployed gateways [8]. Hence, merely capturing the probes from user devices suffices in capturing coarse granularity data for crowd analytics.

The badges on the other hand programmatically send probes at a rate high enough the construct the user trajectories. In short, we used the data from the probes from the badges for behavioral analysis and the data from other devices for crowd analysis.

Figure 1: IEEE 802.11 Probe Request Frame

The format of management frames are shown in Fig. 1. The frame offers information regarding various entities including the type and subtype of the frame, transmitting and receiving devices, MAC address of the AP, sequence number, etc. The transmitting device address (TA) and sequence number fields are used to differentiate each probe from each other. The frame body consists of a number of information elements such as SSID and vendor specific information. SSID information is mandatory in the frame however it can be empty (or null) to specify the probes that are broadcast and not directed towards an AP with a particular SSID.

We next describe the two system components, Wi-Fi badges and gateways that monitor the Wi-Fi channel.

3.2 Wearable Wi-Fi Badge

We begin by describing the hardware and software components of the wearable badge and then briefly discuss its energy footprint.

3.2.1 Hardware

Two concerns mainly drove the hardware design of the badge are its physical size and the functional requirements. Physically the badge should equal the size of a standard credit card (including battery, electronics and all outside connections) to ensure that it could be attached to a standard conference lanyard. Functionally, reprogramming and re-charging of the battery should be achieved by the user. Taking both these concerns into account, we designed a badge that was implemented using a custom designed PCB with a dimension of 85mm x 55mm x 2mm. The badge was composed of an ESP8266[2], a SoC consisting of an ARM-based CPU and a 2.4GHz Wi-Fi controller, a Freescale MMA8452Q accelerometer and associated power regulation and battery charging circuitry. We opted for a ultra thin 180mAH LiPo battery which was regulated

[2]http://www.esp8266.com

to provide the system with power. The accelerometer was used to wake up the system from sleep on motion. Figure 2 shows the schematic of the badge.

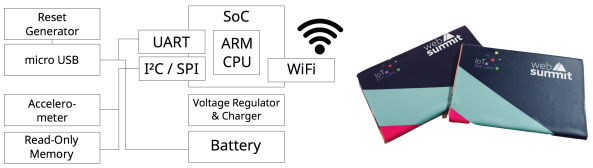

Figure 2: (a) The schematic design and (b) the production version of the badge.

3.2.2 Software

The badge can be charged and programmed using a micro USB connector. We developed an energy aware firmware to broadcast the 802.11 management probes systematically. The firmware also offered remote management functionality using which one could connect to the badge remotely over Wi-Fi and adjust the probe sending rate and power management sensitivity. The firmware implemented a simple algorithm that could detect when the badge was in motion for multiple seconds using the accelerometer and used this as a trigger to send probes. If no motion was detected for a specific period of time, then the badge followed a pre-defined schedule for sending probes (one per minute). The probe functionalities, i.e., custom header construction, emit rate, etc. were implemented in C on top of a modified Wi-Fi stack of ESP firmware. Specifically, we control the `Sequence Control` field of the probe frame to uniquely detect a probe emitted from a specific badge (identified by its MAC address). In addition, we set the `SSID` information element to a particular string. The motion detector was implemented in C using the standard I^2C library.

3.2.3 Energy Management

The ultra thin 180mAH battery had a total capacity to send 1840 probes. Past research has shown that Wi-Fi packet transmission is energy expensive [5]. To ensure that a fully charged badge could operate the whole duration of the events, we used a timer that put the badge into deep sleep after 10 operational hours (08:00 - 18:00), and woke up the badge after 14 hours. During the operational cycle the badge followed the algorithm mentioned earlier for sending probes. This approach ensured that we capture probes for critical times compared to the simpler approach of sending probes at regular intervals until the battery is drained. This design also enabled the badges to operate effectively over the 3 days of the event without the need for recharging.

3.3 Wi-Fi Gateways

In our system, Wi-Fi gateways captured the probes emitted from wearable badges and other smart devices . These gateways were implemented using Raspberry PI 2 Model B [3] one board computers equipped with a Wi-Fi dongle (Ralink 5370) and a programmable LED lamp[4]. The probe capturing functionality was implemented in C++ using the *libtins* library[5]. The LED lamp acted as an indicator that a gateway was operating normally. Figure 3(c) illustrates a gateway and its constituent components.

[3]https://www.raspberrypi.org/products/model-b/
[4]https://blink1.thingm.com
[5]http://libtins.github.io

Once a probe is captured, the gateway looks at the IEEE 802.11 header to extract `TA` and `Sequence number` information to identify each probe. It also checks the `SSID` field to distinguish probes from the badges. The gateway also checks the Radiotap header[6] to get the Received Signal Strength Indication (RSSI). These values along with the local `timestamp` at the gateway are stored in the local storage in the gateway.

As pointed out earlier, the Wi-Fi devices send probes every channel (2.4 GHz or 5GHz band) during their scans. Given the Wi-Fi infrastructure of the event was on the 5GHz band with each device working on a different channel and the badges have only 2.4 GHz Wi-Fi interfaces, we set the Wi-Fi dongle in monitor mode on 2.4GHz band to capture the probes emitted from the badges and smart devices. Since the badges send probes on the less crowded channels, it also increases the probability that these probes are captured.

In addition, we have utilized a daemon in each gateway that monitors the data capturing process. If the process fails due to any reason, the daemon restarts the process but writes the new records to a new file. Hence, we can end up with having a number of data files in a single gateway.

4. SYSTEM DEPLOYMENT

Our study was conducted in early November, 2015 in a three-day industrial exhibition in western Europe. The event attracted over 40K attendees including representatives from major IT companies, small to medium sized startups, media companies, venture capital firms, etc. The $\pm$6000 sq. m. venue was organized strategically to create multiple opportunities for these two groups to meet, interact, and socialize. We deployed 30 gateways in 9 different zones. These zones and the location of the deployed gateways and a snapshot from the exhibition are shown in Figure 3(a,b). Out of the 30 gateways, we were able to retrieve 29 Gateways after the event closing.

5. DATA PRE-PROCESSING

We stored the information of all captured probes in local data files in the Wi-Fi gateways. Since these gateways were isolated and lacked external connectivity, their local times were not synchronized. In total, we obtained 195 files from all 29 gateways. We pre-processed these data files before merging them. Each data entry in the traces consists of `MAC`, `sequence`, `RSSI` and `timestamp` values and a `isBadge` flag that indicates whether the probe is from a badge or not. With pre-processing, we obtain a `global-time` value from the `timestamp` values for each probe.

Because the gateways have no internal battery that can keep the local time up to date, the local time increments from the last known time when they start. Hence, it is assumed that each `timestamp` value from a data file is skewed by a constant value. In order to find this skew, we use received probes as anchors because it is likely that a single probe is received by a number of gateways. Using such commonly received probes, we can find the how far files are apart from each other. Since we recorded the start time of a particular gateway, we could find the difference of the first file it recorded from the actual time. Then, we found the relative difference of each file in all gateways to that particular file.

Each probe is identified by `MAC` and `sequence` fields. However, the `sequence` field is a 12 bit number and it rolls back to zero for each device. A typical device likely emits two probes with the same `sequence` value. However, since the badges are

[6]http://radiotap.org/

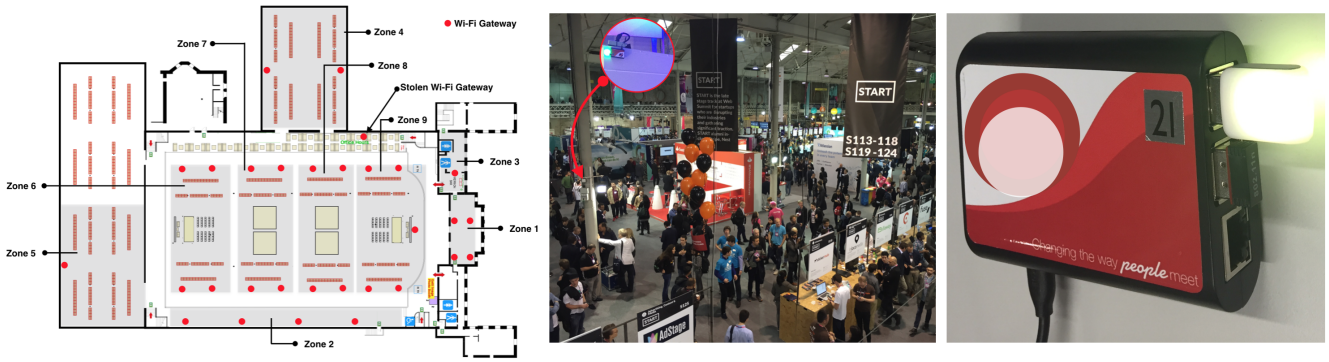

Figure 3: (a) Venue floor plan and deployment map, (b) snapshot of the event and (c) the Wi-Fi gateway.

programmed to send only 1840 probes, every probe from a badge has a unique `MAC` and `sequence` combination. For each file pair, we found the common set of received probes from badges and used them the find the time difference between them. By using the probes from the badges, we were able to merge 122 files. Since the other files did not receive any probes from the badges, we used the entire set of probes they received. For each `MAC`, `sequence` pair in a file that also appears in the merges set, we calculated the difference between the local `timestamp` and the corresponding`global-time`. Since, these pairs are no longer guaranteed to be unique, it is likely to receive multiple values for each pair. We calculated the median difference value and used it as the skew value for that particular file.

In the merged database, we included the `gateway` information to every entry to indicate the particular gateway that captured the probe. Once we merged all the data files, we sorted the data entries by the `global-time` entry.

6. DATA DESCRIPTION

In data, we see that large portion of these records did not contain meaningful information as most of the MAC addresses were seen only a few times or over a very short period of time. 38% of the MAC addresses seen in the traces had only a single probe and more than 85% of the MAC addresses had less than 10 probes in total.

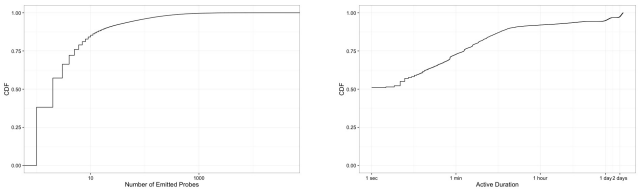

Figure 4: Distribution of the emitted probes and active duration for each device.

As illustrated in the Figure 4 (a), the distribution of probes per device had a very long tail with a small number of non-mobile devices were seen sending over 40K probes. We concur that these were the networked gateways that offered Wi-Fi connectivity, or P2P devices that searched constantly for peers. In addition, about 88% of the MAC addresses were seen only for less than 10 minutes (Figure 4 (b)). We filtered the traces to include MAC addresses that sent at least 300 and at most 1000 probes and were active for at least two different days in the event[7]. Out of the 85 distributed badges, 24 badges either emitted only a few probes, or seen for a very short time, and we also excluded these badges for further analysis.

After this filtering, we ended up with 2526 anonymous device MAC addresses, 61 badges yielding two datasets: a crowd dataset that we have used subsequently for crowd analytics, and a badge dataset that we have used for analyzing the behavior of investors and entrepreneurs. The summary of the entire and the

	# MAC	# Probes	# Badges	# Badge Probes
Total	290740	7054970	85	69191
Filtered	2526	1325364	61	56615

Table 1: Description of the orginal and filtered datasets.

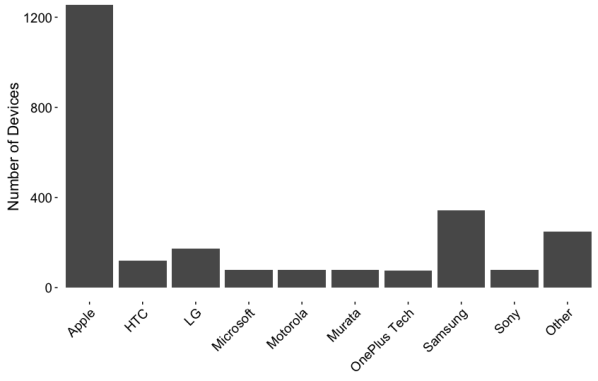

Figure 5: The distribution of device vendors

MAC addresses are structured so that the vendor of the Wi-Fi interface can be determined from the first three octets, which is also known as Organizationally Unique Identifier (OUI). In the filtered

[7]Upper limit on the number is not applied to the badges

traces, we have discovered 67 vendors. Nine most popular vendors have produced 90% of the devices and the emitted probes. The most popular vendor, Apple, accounts for about half of the devices. The distribution of devices per vendor is shown in Figure 5.

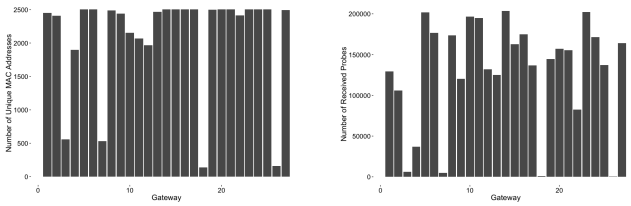

Figure 6: (a) Number of unique MAC addresses discovered by each gateway and (b) the number of probes captured by each gateway

In Figure 6, we show the number of total probes each gateway recorded and the number unique mac addresses that emit these probes. Most gateways receive from almost all 2526 devices apart from a small number of very inactive gateways as a result of the large range of Wi-Fi signals and mobility of users. In average, gateways received more than 129K probes from 2100 devices. Since the range Wi-Fi signals are relatively large and `RSSI` values are highly volatile, it is not straightforward to to get the location of the user from the reception of a probe by a gateway.

We use a proximity ranging mechanism to translate the `RSSI` values associated with a probe from a device that are recorded in a set of gateways to a location relative to the gateways. We formulated this proximity ranging task as a multi-class classification problem by dividing our space in $K = 9$ non-overlapping spatial zones (see Figure 3(a)). We use the Support Vector Machine (SVM) classifier with a One-vs-One scheme [20].

We leave details of the localization mechanism and further analysis on the crowd analytics and behavior traits of the certain groups for future work.

7. CONCLUSION

In this paper, we summarize the elements of a Wi-Fi based monitoring system that aims to collect data in a very large event to study the crowd dynamics in the event and the behavioral analysis of a select group of attendees. We have designed custom Wi-Fi badges and Wi-Fi gateways. We used badges to produce data for behavior analysis. In addition, gateways collect information from other devices present in the event to study the general crowd in the event. We programed the gateways to collect, and the badges to produce only the probe request frames. This enabled us to collect coarse granularity data about the devices present in the event and fine granularity data from the badges that are used of behavior analysis. We also outline pre-processing phase in the system that allowed us to have a coherent data set.

This paper focuses on the system and the data collection phase of our study. We plan to report the lessons we learned from this collected data in future work.

8. REFERENCES

[1] IEEE Standard for Information technology – Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007)*, pages 1–2793, March 2012.

[2] U. G. Acer, A. Boran, C. Forlivesi, W. Liekens, F. Péres-cruz, and F. Kawsar. Sensing WiFi network for personal IoT analytics. In *Internet of Things (IOT), 2015 5th International Conference on the*, pages 104–111, 2015.

[3] U. Blanke, G. Tröster, T. Franke, and P. Lukowicz. Capturing crowd dynamics at large scale events using participatory gps-localization. In *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, pages 1–7, 2014.

[4] T. Bratkovic, B. Antoncic, and M. Ruzzier. Strategic utilization of entrepreneur's resource-based social capital and small firm growth. *Journal of Management & Organization*, 15(04):486–499, 2009.

[5] X. Chen, Y. Chen, M. Dong, and C. Zhang. Demystifying Energy Usage in Smartphones. In *Proceedings of the 51st Annual Design Automation Conference*, pages 70:1–70:5, 2014.

[6] Y. Chon, S. Kim, S. Lee, D. Kim, Y. Kim, and H. Cha. Sensing wifi packets in the air: practicality and implications in urban mobility monitoring. In *Proceedings of the International Joint Conference on Pervasive and Ubiquitous Computing*, pages 189–200, 2014.

[7] T. Elfring and W. Hulsink. Networks in entrepreneurship: The case of high-technology firms. *Small business economics*, 21(4):409–422, 2003.

[8] J. Freudiger. How Talkative is Your Mobile Device?: An Experimental Study of Wi-Fi Probe Requests. In *Proceedings of the 8th ACM Conference on Security & Privacy in Wireless and Mobile Networks*, pages 8:1–8:6, 2015.

[9] R. Friedman, A. Kogan, and Y. Krivolapov. On Power and Throughput Tradeoffs of WiFi and Bluetooth in Smartphones. *IEEE Transactions on Mobile Computing*, 12(7):1363–1376, 2013.

[10] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *Computer*, 34(8):57–66, Aug. 2001.

[11] S. Jamil, A. Basalamah, A. Lbath, and M. Youssef. Hybrid participatory sensing for analyzing group dynamics in the largest annual religious gathering. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 547–558. ACM, 2015.

[12] J. Lampel and A. D. Meyer. Field-configuring events as structuring mechanisms: How conferences, ceremonies, and trade shows constitute new technologies, industries, and markets. *Journal of Management Studies*, 45(6):1025–1035, 2008.

[13] N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, and A. T. Campbell. A survey of mobile phone sensing. *Communications Magazine, IEEE*, 48(9):140–150, 2010.

[14] J. E. Larsen, P. Sapiezynski, A. Stopczynski, M. Mørup, and R. Theodorsen. Crowds, Bluetooth, and Rock'N'Roll: Understanding Music Festival Participant Behavior. In *Proceedings of the 1st ACM International Workshop on Personal Data Meets Distributed Multimedia*, pages 11–18. ACM, 2013.

[15] H. Liu, H. Darabi, P. Banerjee, and J. Liu. Survey of wireless indoor positioning techniques and systems. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 37(6):1067–1080, Nov 2007.

[16] A. Mashhadi, G. Vanderhulst, U. G. Acer, and F. Kawsar. An autonomous reputation framework for physical locations based on wifi signals. In *Proceedings of the 2Nd Workshop*

on Workshop on Physical Analytics, WPA '15, pages 43–46, New York, NY, USA, 2015. ACM.

[17] R. Montoliu and D. Gatica-Perez. Discovering human places of interest from multimodal mobile phone data. In *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia*, 2010.

[18] E. O'Neill, V. Kostakos, T. Kindberg, A. F. gen. Schieck, A. Penn, D. S. Fraser, and T. Jones. Instrumenting the City: Developing Methods for Observing and Understanding the Digital Cityscape. In *UbiComp 2006: 8th International Conference on Ubiquitous Computing*, pages 315–332, 2006.

[19] D. Roggen, M. Wirz, G. Tröster, and D. Helbing. Recognition of crowd behavior from mobile sensors with pattern analysis and graph clustering methods. *arXiv:1109.1664*, 2011.

[20] B. Schölkopf and A. Smola. *Learning with kernels*. M.I.T. Press, 2001.

[21] G. Vanderhulst, A. Mashhadi, M. Dashti, and F. Kawsar. Detecting human encounters from wifi radio signals. In *Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia*, pages 97–108. ACM, 2015.

[22] J. Weppner and P. Lukowicz. Bluetooth based collaborative crowd density estimation with mobile phones. In *Pervasive Computing and Communications (PerCom), 2013 IEEE International Conference on*, pages 193–200, 2013.

[23] G. Zanca, F. Zorzi, A. Zanella, and M. Zorzi. Experimental comparison of rssi-based localization algorithms for indoor wireless sensor networks. In *Proceedings of the workshop on real-world wireless sensor networks*, pages 1–5, 2008.